

CLAIMS

What is claimed is:

1. A system comprising:

- resource data, the resource data including information on a plurality of
5 resources, the resources including a plurality of computers;
application data, the application data including one or more application
profiles, each application profile having a performance profile and a resource
profile, the resource profile including resource utilization information;
policy data, the policy data including one or more application performance
10 policies and one or more resource utilization policies; and
a monitoring agent to monitor the application data for compliance with the
policy data.

2. The system of claim 1, wherein at least one of the applications comprises an
aggregate application executing on at least two of the computers.

3. The system of claim 1, wherein the monitoring agent is further to perform
arbitration within a domain grouping one or more of the computers in response to
a violation of one of the policies.

4. The system of claim 1, wherein the monitoring agent is further to expand a
domain grouping one or more of the computers in response to a policy violation.

5. The system of claim 1, wherein the monitoring agent is further to contract a domain grouping one or more of the computers.

6. The system of claim 1, further comprising domain definition data, the domain definition data including information on a plurality of domains, each domain comprising a grouping of one or more computers, one or more of the domains being a cluster.

7. The system of claim 6, wherein the cluster comprises a first container executing a set of replicated instances of an application on a first set of nodes and a second container having a second set of nodes.

8. The system of claim 7, wherein the monitoring agent is further to transfer a node from the second container to the first container in response to a violation of one of the policies.

9. The system of claim 1, further comprising domain definition data having information on a plurality of domains, each domain comprising a grouping of one or more computers, the domain definition data further including information on the resource utilization of a domain.

10. The system of claim 1, wherein the resource profile further includes resource demand information on the amount of resources an application requires.

11. The system of claim 1, wherein the resource utilization information includes resource consumption information on the amount of resources an application is currently assigned.

12. The system of claim 1, wherein the resource utilization information includes at least one of resource consumption information on the amount of resources an application is currently using, and resource consumption information on the amount of resources an application has used over a period of time.

13. The system of claim 1, wherein one of the computers is associated with a container to execute one of the applications.

14. The system of claim 1, wherein one of the computers is associated with a plurality of containers, each container to execute one of the applications.

15. The system of claim 14, wherein the policy data further includes one or more container utilization policies, each utilization policy associated with one of the containers.

16. The system of claim 14, wherein at least one of the containers is a partition.

17. The system of claim 16, wherein the monitoring agent is further to resize the partition in response to a violation of one of the policies.

18. The system of claim 16, wherein the partition is a hardware partition.
19. The system of claim 16, wherein the partition is a software-based partition.
20. The system of claim 14, wherein at least one of the containers is a processor set.
21. The system of claim 14, wherein at least one of the containers is a sub-CPU resource partition.
22. The system of claim 1, wherein the performance information includes response time.
23. The system of claim 1, wherein one or more of the application profiles includes resource allocation information for the associated application.
24. The system of claim 1, wherein one or more of the application profiles further includes instructions for installing the associated application.
25. The system of claim 22, wherein the instructions further include instructions for configuring the associated application.

26. The system of claim 1, wherein the policy data further includes one or more resource utilization policies.

27. The system of claim 1, wherein the performance policies have a relative associated priority.

28. A method comprising:

monitoring application data for compliance with one or more performance policies, the application data including one or more application profiles, each application profile having a performance profile and a resource profile, the
5 resource profile including resource utilization information associated with an application, each application executing in a container associated with a domain, each domain including one or more resources, the resources including a plurality of computers;

in response to a policy violation, automatically enforcing the policy by
10 expanding a first one of the containers.

29. The method of claim 28, wherein the first container comprises a partition and expanding the first container comprises resizing the partition.

30. The method of claim 28, wherein the domain associated with the first container comprises a cluster.

31. The method of claim 30, wherein expanding the first container comprises transferring a node associated with a second container, the second container being in the domain associated with the first container, to the first container.

32. The method of claim 28, further comprising in response to a second policy violation, providing a message to a user.

33. The method of claim 32, wherein the message comprises a message that a lower priority policy cannot be met.